

Composite likelihood inference for spatio-temporal data on multicolor cell growth

Puxue Qiao, Christina Mølck, Davide Ferrari,
Cameron J. Nowell, and Frédéric Hollande

October 19, 2016

Abstract

Multicolor cell spatio-temporal image data have become important to investigate organ regeneration, malignant outgrowth or immune responses by tracking different cell types both in vivo and in vitro. Statistical modeling of image data from common longitudinal cell experiments poses significant challenges due to the presence of complex spatio-temporal interactions between different cell types and difficulties related to measurement of single cell trajectories. Current analysis methods oversimplify the problem to the benefit of computational feasibility, often not providing a full statistical treatment of the spatio-temporal effects affecting cell growth. In this paper, we propose a conditional spatial autoregressive model to describe multivariate count cell data and develop a new parameter estimator within the composite likelihood inference framework. The proposed methodology is computationally tractable and enables researchers to estimate a complete statistical model of multicolor cell growth. The proposed methods are supported by real experimental data where we investigate how interactions between cancer cells and fibroblasts, which are normally present in the tumor microenvironment, affect their growth.

Keywords: Spatio-temporal lattice model, count data, composite likelihood estimation, multicolor cell growth

1 Introduction

Longitudinal image data based on fluorescent proteins play a crucial role for both in vivo and in vitro analysis of various biological processes such as gene expression and cell lineage fate. Assessing the growth patterns of different cell types within a heterogeneous population and monitoring their spatio-temporal interactions enables biomedical researchers to determine the role of different cell types in important processes such as organ regeneration, malignant outgrowth or immune responses under various experimental conditions. Our work is motivated by longitudinal image data collected from multicolor live-cell imaging growth experiments. Using a high-content imaging system, we were able to acquire characteristics for each individual cell at subsequent times, including fluorescent properties, spatial coordinates, and morphological features.

In longitudinal growth experiments, the two important goals are to determine growth rates for different cell populations and to assess how spatio-temporal interactions between cell types may affect their growth. Whilst a wide range of descriptive data analysis approaches have been used in applications, inference based on a comprehensive model of multicolor cell data is an open research area. The main challenges are related to the presence of complicated spatio-temporal interactions amongst cells and difficulties related to tracking individual cells across time from image data. Typical longitudinal experiments consist of a relatively small number of measurements (e.g. 5 to 20 images taken every few hours), which is adequate for monitoring cell growth. Tracking individual cells would typically require more frequent measurements, complicating the practicality of the experiments in terms of the storage cost of very large image files and the toxicity on the cells induced by the imaging process.

Although tracking individual cell trajectories is difficult due to cell migration, overlapping cells, changes in cell morphology, image artifacts, cell death and division, obtaining cell counts by cell type (represented by a certain color) is straightforward and can be easily automated. To describe the spatial distribution for different cell types, we propose to divide an image into a number of contiguous regions (tiles) to form a regular lattice as shown in Figure 1 (a). We then record and model the frequency of cells of different colors in each tile

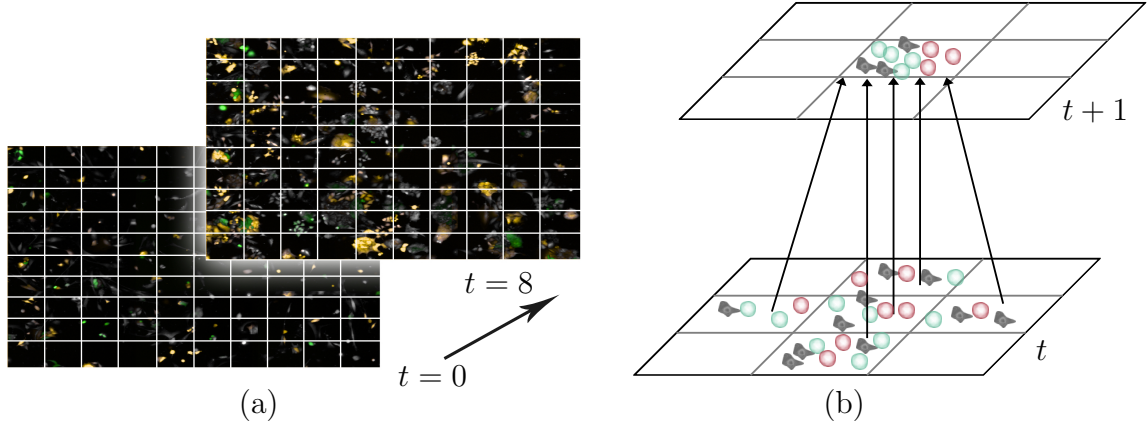


Figure 1: (a) Microscope images for the cancer cell growth data obtained from a high-content imager (Operetta, Perkin Elmer) at the initial and final time points of the experiment. In each image, colors for non-fluorescent fibroblasts, as well as red and green fluorescent cancer cells are merged. (b) Illustration of the local structure for the conditional autoregressive spatio-temporal Model in (1) and (2). The two planes correspond to 3×3 tiles at times t and $t + 1$. The average number of cells of color c in a given tile at time $t + 1$ is assumed to depend on the number of cells of other colors in contiguous neighboring tiles at time t .

at subsequent time points.

The first contribution of this paper is the development of a conditional spatial autoregressive model for multivariate count data on tiled images and its application in the context of longitudinal cancer cell monitoring experiments. Our model enables us to measure the effect on the growth rate of each cell population and changes due to local cross-population interactions. Specifically, we consider a multivariate Poisson model with intensity modeled as a nonlinear function of spatio-temporal interactions between cells belonging to different color groups in neighboring tiles as illustrated in Figure 1 (b). The main advantage of the proposed framework is that it enables one to accommodate spatio-temporal cell interactions for heterogeneous cell populations within a relatively parsimonious statistical model.

Spatial lattice models are well established tools and have been applied in numerous areas of empirical research; see, e.g., Cressie (2015), Chapter 6. However, the literature of spatio-temporal lattice models is not as developed as that of spatial lattice models. Gelfand et al. (2003) propose a linear model with spatially correlated responses and consider no space-time interactions, spatial dependence attenuated across time, and spatially varying parameters. Bowman (2007) develop mixed effect spatio-temporal linear models to analyze functional neuroimaging data. Kang et al. (2012) propose a spatio-spectral mixed-effect model to analyze functional magnetic resonance imaging (fMRI) data consisting of a set

of discrete time series measured on three-dimensional volume elements called voxels. Their model takes into account three types of correlations: local spatial correlation between voxels, global spatial correlation, and the temporal correlation within a voxel. The literature on nonlinear spatio-temporal models includes a growing number of heterogeneous applications. For example, Zhu et al. (2008) discuss a spatio-temporal autologistic regression model with a binary response. Zheng and Zhu (2012) extend the Markov random field (MRF) model to include temporal terms and multiple-response variables. Overall, statistical modelling of nonlinear and non-Gaussian spatio-temporal dynamical processes is a relatively young area and there is much to be done from both modeling and estimation viewpoints.

The second contribution of the present paper is the development of computationally efficient inference and model selection tools for the proposed discrete spatio-temporal model. Maximum likelihood (ML) estimation is known to be the best option in terms of accuracy for parametric models. However, the full likelihood function is unavailable in our setting due to the computational complexity associated with the full multivariate Poisson model. To overcome this issue, we develop an alternative inference approach within the composite likelihood (CL) framework. In CL estimation, difficulties related to evaluation of the full likelihood functions are avoided by considering an approximation obtained by merging a number of smaller likelihood objects. Besag (1974) was an early proponent of CL methods for spatial data and Lindsay (1988) developed composite likelihood inference in its generality. Over the years, CL methods have proved useful in a range of complex applications. See Varin et al. (2011) for an overview of the methods and applications. Although in CL estimation some statistical efficiency is lost compared to ML, the very low computational burden makes CL methods appropriate when working with the complex spatio-temporal data motivating this paper.

There is a growing literature describing implementations of CL methods in the context of spatio-temporal models. Among others, Zhu et al. (2005) use CL estimation in a spatio-temporal regression model on lattice with binary response. Hyrien et al. (2005) and Hyrien (2007) analyze cell growth across time using a CL parametric estimation, with conditional distributions of cell counts approximated by simulation. Chen et al. (2011) use CL estimation to analyze longitudinal data on multiple cell clones with marginal distributions on the

number of cells approximated by simulation.

Since the model complexity can be potentially very large in the presence of many cell types, it is also important to address the question of how to select an appropriate model by retaining only the meaningful spatio-temporal interactions between cell populations. However, common model selection criteria for parametric models, such as the Akaike and the Bayesian information criteria (AIC and BIC), cannot be used in our setting since they rely on the availability of the full model and its correct specification. To overcome these issues, we propose a model selection strategy by adapting the composite likelihood information criterion of Varin and Vidoni (2005) in the context of our spatio-temporal model.

The remainder of the paper is organized as follows. In Section 2, we introduce the conditional spatio-temporal lattice model for multivariate count data and develop composite likelihood inference tools. In the same section, we describe the CL information criterion for model selection. In Section 3, we study the performance of the new estimator using simulated data. In Section 4, we apply our method to analyze a dataset from an in-vitro experiment on cancer cells co-cultured with fibroblasts. In Section 5, we conclude and give final remarks.

2 Methods

2.1 Multicolor autoregressive model on the lattice

Let \mathcal{L} be a discrete regular lattice and use \sim to denote a symmetric irreflexive relation on \mathcal{L} . In the context of our application, the lattice is obtained by tiling a microscope image into $n \times n$ rectangular tiles. An example of a tiled image with $n = 10$ is shown in Figure 1 (a). Here we assume a square lattice where the pair of tiles $\{i, j\}$ are neighbors with $i \sim j$ if tiles i and j share the same border or coincide ($i = j$), and use $n_{\mathcal{L}} = |\mathcal{L}| = n \times n$ to denote the total number of tiles. Each tile may contain cells of different colors; thus, we let $\mathcal{C} = \{1, \dots, n_{\mathcal{C}}\}$ be a finite set of colors and denote by $n_{\mathcal{C}}$ the total number of colors. Let $Y = \{Y_t^{(c)}, c \in \mathcal{C}; t \leq T\}$ be the sample of observations with $Y_t^{(c)} = (Y_{1,t}^{(c)}, \dots, Y_{n_{\mathcal{L}},t}^{(c)})^\top$ being the vector of observed frequencies for color c on the lattice \mathcal{L} at time t . The joint distribution for the spatio-temporal process on the lattice is difficult to specify, due to local spatial interactions for neighboring tiles and global interactions occurring at the level of the

entire image. An additional issue is that cells tend to be clustered together due to the cell division process and other biological mechanisms; thus it is common to observe low or zero counts in a considerable portion of tiles.

We suppose $Y_t^{(c)}$ follows the $n_{\mathcal{L}}$ -variate multivariate Poisson distribution $Y_t^{(c)} \sim \text{MultPois}(\lambda_t^{(c)})$, where $\lambda_t^{(c)} = (\lambda_{1,t}^{(c)}, \dots, \lambda_{n_{\mathcal{L}},t}^{(c)})^\top$ denotes the $n_{\mathcal{L}} \times 1$ mean vector. In general, a k -variate multivariate Poisson distribution is defined as a distribution on \mathbb{N}^k with univariate Poisson marginals. Here we assume that the count for the i th tile $Y_{i,t}^{(c)}$ follows a marginal Poisson distribution $Y_{i,t}^{(c)} \sim \text{Pois}(\lambda_{i,t}^{(c)})$, with intensity modeled by the canonical log-link $v_{i,t}^{(c)} = \log \lambda_{i,t}^{(c)}$, where $v_{i,t}^{(c)}$ takes the following spatial autoregressive form:

$$v_{i,t}^{(c)} = \alpha^{(c)} + \sum_{c' \in \mathcal{C}} \beta^{(c|c')} S_{i,t-1}^{(c')}, \quad (1)$$

$$S_{i,t-1}^{(c')} = \frac{1}{n_i} \sum_{i \sim j: j \in \mathcal{L}} \log(1 + Y_{j,t-1}^{(c')}), \quad (2)$$

for all $c \in \mathcal{C}, 0 \leq t \leq T$, with $n_i = \{\#j : i \sim j, j \in \mathcal{L}\}$ being the number of tiles in a neighborhood of tile i .

The elements of the parameter vector $\alpha = (\alpha^{(1)}, \dots, \alpha^{(n_{\mathcal{C}})})^\top$ are main effects corresponding to a baseline average count for cells of different colors. The spatio-temporal interactions are measured by the statistic $S_{i,t-1}^{(c')}$ in (2), which essentially counts the number of cells of color c' in the neighborhood of tile i at time $t - 1$. Hence, the autoregressive parameter $\beta^{(c|c')}$ is interpreted as positive or negative change in the average number of cells with color c , due to interactions with cells of color c' . A positive (or a negative) sign of $\beta^{(c|c')}$ means that the presence of cells of color c' in neighboring tiles promotes (or inhibits) the growth of cells of color c . The spatio-temporal effects $\beta^{c|c'}, c, c' \in \mathcal{C}$, are collected in the $n_{\mathcal{C}} \times n_{\mathcal{C}}$ weighted incidence matrix \mathcal{B} . This may be used to generate weighted directed graphs, as shown in the example of Figure 2, where the nodes of the directed graph correspond to cell types, and the directed edges are negative or positive spatio-temporal interactions between cell types.

Similar to the univariate autoregressive poisson model of Fokianos and Tjøstheim (2011), the log-linear form for the intensity in Equation (1) offers several advantages compared to

the more commonly used linear form. First, note that lagged observations are included via the mapping $\log(1 + Y_{i,t-1}^{(c)})$; this can handle zero counts, while mapping zeros of $Y_{i,t-1}^{(c)}$ into $\log(Y_{i,t-1}^{(c)})$ may generate infinite intensities. Second, $\lambda_{i,t}^{(c)}$ and $Y_{i,t-1}^{(c)}$ are transformed on the same scale; in our experience, this renders inference quite stable. Finally, this model can accommodate both positive and negative correlations, which cannot be achieved if we included $Y_{i,t-1}^{(c)}$ instead of $\log(Y_{i,t-1}^{(c)} + 1)$. For example, with the model $v_{i,t} = \alpha + \beta Y_{i,t-1}$ for a single color, the intensity would be $\lambda_{i,t} = \exp(\alpha) \exp(\beta Y_{i,t-1})$, which may lead to instability of the Poisson means if $\beta > 0$ since $\lambda_{i,t}$ would be allowed to increase exponentially fast.

2.2 Composite likelihood inference

Let θ be the overall parameter vector $\theta = (\alpha^\top, \text{vec}(\mathcal{B})^\top)^\top \in \mathbb{R}^p$, with α and \mathcal{B} defined in Section 2.1 and $p = n_{\mathcal{C}}(1 + n_{\mathcal{C}})$ being the total number of parameters. In this section, we develop inference based on a composite likelihood approach. The full likelihood function can be written as

$$L(\theta; Y) = \prod_{1 \leq t \leq T} \prod_{c \in \mathcal{C}} P(Y_t^{(c)} | Y_{t-1}^{(1)}, \dots, Y_{t-1}^{(n_{\mathcal{C}})}; \theta),$$

where the conditional distribution of $Y_t^{(c)}$ given $Y_{t-1}^{(1)}, \dots, Y_{t-1}^{(n_{\mathcal{C}})}$ is assumed to be a multivariate Poisson $\text{MultPois}(\lambda_t^{(1)}, \dots, \lambda_t^{(n_{\mathcal{C}})})$. Following Karlis and Meligkotsidou (2005), a multivariate Poisson distribution is defined through a k -vector Z with elements being independent Poisson random variables and a mapping $g : \mathbb{N}^k \mapsto \mathbb{N}^k$. Specifically, the probability mass function of $X = g(z)$ is

$$P(X = x) = \sum_{z \in g^{-1}(x)} P(Z = z), \quad (3)$$

where $g^{-1}(x) \subseteq \mathbb{N}^k$ is the inverse image of x under g . A main issue limiting the direct use of the multivariate Poisson model in applications is the computational complexity of the sum in (3) which involves $\mathcal{O}(2^k)$ terms (Karlis and Meligkotsidou, 2005). For example, evaluating the distribution function of an image with $n_{\mathcal{C}} = 100$ tiles under a general covariance structure would require summing approximately 1.3×10^{30} terms in (3).

To avoid difficulties related to evaluation of the full likelihood $L(\theta; Y)$, we propose an

alternative objective function, called composite likelihood (CL) function, by merging simpler conditional univariate Poisson likelihoods. The composite likelihood (CL) function is

$$CL(\theta) = \prod_{1 \leq t \leq T} \prod_{c \in \mathcal{C}} \prod_{i \in \mathcal{L}} P(Y_{i,t}^{(c)} | Y_{t-1}; \theta)^{w_i} = \prod_{1 \leq t \leq T} \prod_{c \in \mathcal{C}} \prod_{i \in \mathcal{L}} \left(e^{-\lambda_{i,t}^{(c)}(\theta)} \frac{\lambda_{i,t}^{(c)}(\theta)^{y_{i,t}^{(c)}}}{y_{i,t}^{(c)}!} \right)^{w_i}, \quad (4)$$

where $\lambda_{i,t}^{(c)}(\theta)$ is the expected number of cells with color c in tile i at time t , defined in (1), and $w = (w_1, \dots, w_{n_{\mathcal{L}}})^\top$ are given constants. The maximum composite likelihood estimator (MCLE), $\hat{\theta}$, is obtained by maximizing the log-composite likelihood function

$$cl(\theta) = \sum_{i \in \mathcal{L}} w_i \ell_i(\theta) \equiv \sum_{i \in \mathcal{L}} \sum_{1 \leq t \leq T} \sum_{c \in \mathcal{C}} w_i \left\{ Y_{i,t}^{(c)} v_{i,t}^{(c)}(\theta) - \exp \left(v_{i,t}^{(c)}(\theta) \right) \right\}, \quad (5)$$

where $v_{i,t}^{(c)}(\theta) \equiv \log \lambda_{i,t}^{(c)}(\theta)$. Equivalently, $\hat{\theta}$ is formed by solving the weighted estimating equations

$$0 = \nabla cl(\theta) = \sum_{i \in \mathcal{L}} \sum_{1 \leq t \leq T} w_i u_{i,t}(\theta) \equiv \sum_{i \in \mathcal{L}} \sum_{1 \leq t \leq T} w_i \left\{ Y_{i,t}^{(c)} - \exp \left(v_{i,t}^{(c)}(\theta) \right) \right\} \nabla v_{i,t}^{(c)}(\theta), \quad (6)$$

where ∇ denotes the gradient operator with respect to θ and $\nabla v_{i,t}^{(c)}(\theta) = (1, S_{i,t-1}^{(1)}, \dots, S_{i,t-1}^{(n_c)})^\top$.

For computational simplicity, we take uniform weights $w = (1, \dots, 1)^\top$, which is a common choice in the literature of CL estimation (Varin et al., 2011). Our empirical results also show that uniform weights perform reasonably well in terms of estimation accuracy in all our numerical examples. The solution to Equation (6) is obtained by a standard Fisher scoring algorithm, which is found to be stable and converges fast in all our numerical examples.

2.3 Standard error evaluation

As the data dimension $n_{\mathcal{L}} \times T$ increases, theory of unbiased estimating equations (e.g, see Heyde (1987)) suggests that under appropriate regularity conditions, the estimator $\hat{\theta}$ follows approximately a p -variate normal distribution with mean equal to the true parameter θ_0 and asymptotic variance $V(\hat{\theta}) = G^{-1}(\theta_0)$, where $G(\theta) \equiv H^{-1}(\theta)K(\theta)H^{-1}(\theta)$ is the $p \times p$ Godambe information matrix associated with the CL Equation (6), and $H(\theta)$ and $K(\theta)$ are

the $p \times p$ matrices

$$H(\theta) = -E [\nabla^2 c\ell(\theta)] = -E \left(\sum_{i \in \mathcal{L}} w_i \nabla u_i(\theta) \right), \quad (7)$$

and

$$K(\theta) = \text{Var}[\nabla c\ell(\theta)] = \sum_{i \in \mathcal{L}} w_i^2 \text{Var}(u_i(\theta)) + 2 \sum_{i \neq j: i, j \in \mathcal{L}} w_i w_j \text{Cov}(u_i(\theta), u_j(\theta)), \quad (8)$$

with $u_i(\theta) = u_{i,1}(\theta) + \dots + u_{i,T}(\theta)$ being the partial score function for the i th tile.

The main regularity assumptions required to guarantee a proper asymptotic behavior concern existence of the expectations in the expressions of $H(\theta)$ and $K(\theta)$ and, importantly, the order of the second term in $K(\theta)$ involving all pairwise covariances. For example, if $w_i > 0$ and $\text{Cov}(u_i(\theta), u_j(\theta)) > 0$ in a Loewner order sense for all $i, j \in \mathcal{L}$, then $\det G(\theta)$ will approach zero as $n_{\mathcal{L}}$ increases. In the special case where the true model is given by independent Poissons on the lattice, the CL function coincides with the full likelihood. Thus, we have $K(\theta) = H(\theta)$, since $\text{Cov}(u_i(\theta), u_j(\theta)) = 0$ for all $i, j \in \mathcal{L}$, and the Godambe information matrix $G(\theta) = H^{-1}(\theta)$ coincides with the Fisher information matrix.

Direct evaluation of the matrices $H(\theta)$ and $K(\theta)$ is challenging since the expectations in (7) and (8) depend on the full model, which is intractable. Therefore, similarly to Fokianos and Kedem (2004), we consider estimating $V(\theta)$ by the sandwich estimator

$$\hat{V}_{\text{snd}}(\hat{\theta}) = \hat{H}^{-1}(\hat{\theta}) \hat{K}(\hat{\theta}) \hat{H}^{-1}(\hat{\theta}), \quad (9)$$

where $\hat{\theta}$ is the MCLE defined in (5), whilst \hat{H} and \hat{K} are estimators of H and K defined as

$$\hat{H}(\theta) = \sum_{i \in \mathcal{L}} \sum_{t \leq T} \sum_{c \in \mathcal{C}} w_i \exp \left[v_{i,t}^{(c)}(\theta) \right] \left[\nabla v_{i,t}^{(c)}(\theta) \right] \left[\nabla v_{i,t}^{(c)}(\theta) \right]^{\top}, \quad (10)$$

$$\begin{aligned} \hat{K}(\theta) &= \sum_{i, j \in \mathcal{L}} \sum_{t \leq T} \sum_{c \in \mathcal{C}} w_i w_j \left\{ Y_{i,t}^{(c)} - \exp \left(v_{i,t}^{(c)}(\theta) \right) \right\} \left\{ Y_{j,t}^{(c)} - \exp \left(v_{j,t}^{(c)}(\theta) \right) \right\} \\ &\quad \times \left[\nabla v_{i,t}^{(c)}(\theta) \right] \left[\nabla v_{j,t}^{(c)}(\theta) \right]^{\top}. \end{aligned} \quad (11)$$

Note that the above estimators approximate the quantities in formulas (7) and (8) by con-

ditional expectations. Our numerical results suggest that the above variance approximation yields confidence intervals with coverage close to the nominal level $(1 - \alpha)$ when the total number of time points, T , is relatively large. However, when T is small ($T < 25$), as it is the case in our real-data application, the standard errors may be too optimistic with coverage probability smaller than $(1 - \alpha)$.

Besides the above formulas, we also consider confidence intervals obtained by a parametric bootstrap approach. Specifically, we generate B bootstrap samples $Y_{(1)}^*, \dots, Y_{(B)}^*$ by sampling at subsequent times from the conditional model specified in Equations (1) and (2) with $\theta = \hat{\theta}$. From such bootstrap samples, we obtain corresponding MCLEs, $\hat{\theta}_{(1)}^*, \dots, \hat{\theta}_{(B)}^*$, which are used to estimate $V(\theta_0)$ by the usual covariance estimator $\hat{V}_{boot}(\hat{\theta}) = \sum_{b=1}^B (\hat{\theta}_{(b)}^* - \bar{\theta}^*)^2 / (B - 1)$, where $\bar{\theta}^* = \sum_{b=1}^B \hat{\theta}_{(b)}^* / B$.

Finally, a $(1 - \alpha)100\%$ confidence interval for θ_j is obtained as $\hat{\theta}_j \pm z_{1-\alpha/2} \{\hat{V}\}_{jj}^{1/2}$, where z_q is the q -quantile of a standard normal distribution, and \hat{V} is an estimate of $V(\theta_0)$ obtained by either Equation (9) or the bootstrap sampling method.

2.4 Model selection

Model selection plays an important role by balancing goodness-of-fit and model complexity. Although traditional model selection approaches such as the Akaike or the Bayesian information criteria (AIC and BIC) are well-established methods in the context of maximum likelihood estimation, they are not appropriate when the likelihood function is mis-specified, as is the case for the objective function $c\ell(\theta)$ defined in Equation (5); see Chapter 2.8, Claeskens et al. (2008).

To overcome this issue, we use the composite likelihood information criterion discussed in Varin and Vidoni (2005), which is specifically designed for composite likelihood functions. Following Varin and Vidoni (2005), model selection can be performed by minimising the weighted sum of Kullback-Leibler (KL) divergences

$$KL(\theta; \theta_0) = \sum_{i \in \mathcal{L}} w_i E_{\theta_0} [\ell_i(\theta_0) - \ell_i(\theta)],$$

where θ_0 denotes the true parameter, $\ell_i(\theta)$ is the partial likelihood function on tile i defined in

Equation (5), and $w = (w_1, \dots, w_{n_{\mathcal{L}}})^\top$ is such that $KL(\theta; \theta_0) \geq 0$. The maximized composite log-likelihood $c\ell(\hat{\theta})$ is a biased estimator of $KL(\theta; \theta_0)$ with bias of order $\text{tr}(K(\tilde{\theta})H(\tilde{\theta})^{-1})$, where $\tilde{\theta}$ is the minimizer of $KL(\theta_0, \theta)$ and K and H are $p \times p$ matrices defined in (7) and (8). Then, we select the model minimizing the CL information criterion $\text{CLIC} = -2c\ell(\hat{\theta}) + 2\text{tr}(K(\hat{\theta})H^{-1}(\hat{\theta}))$. As an alternative, we also consider the CL Bayesian information criterion $\text{CL-BIC} = -2c\ell(\hat{\theta}) + \log(n_{\mathcal{L}}T)\text{tr}(K(\hat{\theta})H^{-1}(\hat{\theta}))$ similarly to Gao and Song (2010).

In our numerical examples, we estimate $K(\tilde{\theta})H(\tilde{\theta})^{-1}$ in the penalty term of CLIC and CL-BIC by $\hat{H}(\hat{\theta})\hat{V}_{boot}(\hat{\theta})$, since the latter is shown to be more stable in our numerical experiments compared to direct estimator of $K(\theta)$ by Formula (11).

3 Monte Carlo simulations

In our Monte Carlo experiments, we generate data from a Multivariate Poisson model as follows. At time $t = 0$, we populate $n_{\mathcal{L}}$ tiles using equal counts for cells of different colors. For $t = 1, \dots, T$, observations are drawn from the multivariate Poisson model

$$Y_t^{(c)} | Y_{t-1}^{(1)}, \dots, Y_{t-1}^{(n_{\mathcal{L}})} \sim \text{MultPoisson}(\lambda_t^{(c)}, \Sigma(\rho)), \quad c \in \mathcal{C}, \quad (12)$$

where $\lambda_t^{(c)} = (\lambda_{1,t}^{(c)}, \dots, \lambda_{n_{\mathcal{L}},t}^{(c)})^\top$ is a vector of color-specific rates for the $n_{\mathcal{L}}$ tiles and $\Sigma(\rho)$ is a $n_{\mathcal{L}} \times n_{\mathcal{L}}$ correlation matrix with all off-diagonal elements equal to $0 \leq \rho < 1$. In our simulations, the correlation matrix $\Sigma(\rho)$ expresses the residual global interactions between tiles that is not captured by the local conditional model (1). Random draws from (12) are obtained by the function `genPoisNor` in the R package `PoisNor` (Amatya and Demirtas, 2015).

The rate vector $\lambda_t^{(c)}$ defined in Section 2.1 contains coefficients $\beta^{(c|c')}$ defined as elements of the weighted incidence matrix

$$\mathcal{B} = \begin{pmatrix} 1 & -1 & 1 \\ 1 & 1 & -1 \\ -1 & 1 & 1 \end{pmatrix} \otimes A,$$

where \otimes denotes element-wise matrix multiplication and A is a $n_{\mathcal{C}} \times n_{\mathcal{C}}$ constant matrix of weights. To assess the performance of MCLE under different settings concerning the size and sparsity of autoregressive coefficients in $\lambda_{i,t}^{(c)}$, we consider the following three models with different choices of A .

- Model 1: $\{A\}_{c,c'} = 0.7$, for all $c, c' = 1, \dots, n_{\mathcal{C}}$; all the effects in \mathcal{B} have the same size;
- Model 2: $\text{Vec}(A) = (0.05, 0.15, \dots, 0.85)^\top$, the effects forming \mathcal{B} have decreasing sizes;
- Model 3: $\{A\}_{c,c'} = 0$, if $c \neq 1$ and $c' \neq c$, and 0.7 otherwise. This model is similar to Model 1, but some interactions are exactly zero.

We set $\alpha^{(1)} = \dots = \alpha^{(n_{\mathcal{C}})} = -0.1$ for all three models. The above parameter choices reflect the situation where the generated process Y has a moderate growth.

In Tables 1 and 2, we show results based on 1000 samples generated from Models 1-3, for $n = 25, n_{\mathcal{C}} = 3$ and $T = 10, 25, 100$. In Table 1, we show the Monte Carlo estimates of squared bias ($\widehat{\text{Bias}}^2 = (\hat{E}(\hat{\theta}) - \theta_0)^2$) and estimates of variance ($\widehat{\text{Var}} = \hat{E}(\hat{\theta} - E(\hat{\theta}))^2$). Both squared bias and variance of our estimator are quite small in all three models, and decrease as T gets larger. The variances of Model 2 are slightly larger than those in the other two models due to the increasing difficulty in estimating parameters close to zero.

	$T = 10$		$T = 25$		$T = 100$	
	$\widehat{\text{Bias}}^2$	$\widehat{\text{Var}}$	$\widehat{\text{Bias}}^2$	$\widehat{\text{Var}}$	$\widehat{\text{Bias}}^2$	$\widehat{\text{Var}}$
Model 1	0.89 (0.84)	3.37 (0.15)	0.54 (0.48)	1.47 (0.06)	0.14 (0.12)	0.45 (0.02)
Model 2	2.08 (1.91)	5.84 (0.26)	1.32 (0.97)	2.46 (0.11)	0.38 (0.46)	2.04 (0.15)
Model 3	3.57 (1.70)	4.07 (0.18)	2.40 (1.30)	2.98 (0.15)	0.20 (0.15)	0.54 (0.02)

Table 1: Monte Carlo estimates for squared bias ($\times 10^{-5}$) and variance ($\times 10^{-3}$) of the MCLE for three models with time points $T = 10, 25, 100$. Estimated standard errors are shown in parenthesis. The three models differ in terms of the coefficients $\beta^{cl'}$, $c, c' \in \mathcal{C}$ as described in Section 3: Non-zero equal effects (Model 1), non-zero decreasing interactions (Model 2), and sparse effects (Model 3). For all models, $\alpha_c = -0.1$, $c = 1, 2, 3$. Estimates are based on 1000 Monte Carlo runs.

In Table 2, we report the coverage probability for symmetric confidence intervals with the usual form $\hat{\theta} \pm z_{1-\alpha/2} \widehat{sd}(\hat{\theta})$, where z_q is the q -quantile for a standard normal distribution, with $\alpha = 0.01, 0.05, 0.10$. The standard error, $\widehat{sd}(\hat{\theta})$, is obtained by the sandwich and

the parametric bootstrap estimate, \hat{V}_{snd} and \hat{V}_{boot} , described in Section 2.3. The coverage probability of the bootstrap confidence intervals are very close to the nominal level. Although the coverage of confidence intervals based on the sandwich estimator (9) tends to underestimate $1 - \alpha$, but tends to converge to it as T increases.

		$T = 10$		$T = 25$		$T = 100$	
		\hat{V}_{boot}	\hat{V}_{snd}	\hat{V}_{boot}	\hat{V}_{snd}	\hat{V}_{boot}	\hat{V}_{snd}
$\alpha = 0.01$	Model 1	98.6	94.6	98.9	96.4	98.7	98.0
	Model 2	99.1	93.8	98.9	96.5	98.7	98.1
	Model 3	98.8	95.0	98.6	98.4	98.9	98.0
$\alpha = 0.05$	Model 1	94.3	87.2	95.0	89.9	94.7	92.4
	Model 2	95.1	86.3	95.5	89.8	94.3	93.8
	Model 3	94.5	88.0	94.5	95.0	94.5	92.3
$\alpha = 0.10$	Model 1	89.7	80.8	90.3	83.9	89.9	86.8
	Model 2	90.1	79.4	90.8	83.4	89.5	89.3
	Model 3	89.7	82.0	89.5	91.7	89.6	86.9

Table 2: Monte Carlo estimates for the coverage probability of $(1 - \alpha)\%$ confidence intervals $\hat{\theta} \pm z_{1-\alpha/2} \hat{sd}(\hat{\theta})$, with $\hat{sd}(\hat{\theta})$ obtained using bootstrap (\hat{V}_{boot}) and sandwich (\hat{V}_{snd}) estimators in Section 2 and 3. The three models differ in terms of the coefficients $\beta^{cl'}$, $c, c' \in \mathcal{C}$ as described in Section 3: Non-zero equal effects (Model 1), non-zero decreasing interactions (Model 2), and sparse effects (Model 3). For all models, $\alpha_c = -0.1$, $c = 1, 2, 3$, estimates are based on 1000 Monte Carlo runs.

In Table 3, we show results for the model selection based on 500 Monte Carlo samples from Model 3 using CLIC and CL-BIC criteria for $n = 10$ and $T = 10, 25$. We report Type I error (a term is not selected when it actually belongs to the true model), Type II error (a term is selected when it is not in the true model) and

$$F - \text{statistic} = 2 \frac{|\text{supp}(\hat{\beta}) \cap \text{supp}(\beta)|}{|\text{supp}(\hat{\beta})| + |\text{supp}(\beta)|}, \quad (13)$$

where $\text{supp}(\beta) = \{(c, c') : \beta^{cl'} > 0\}$. Model selection with both criteria performs better with larger T . In general, CL-BIC outperforms CLIC, particularly when $T = 25$ the F-statistic for CL-BIC is very close to one, meaning that there is a very high chance that the method selects the correct model.

	$T = 10$			$T = 25$		
	Type I	Type II	F -Stat	Type I	Type II	F -Stat
CLIC	0.03	7.60	89.72	0.00	7.20	90.25
CL-BIC	0.43	0.67	98.44	0.00	0.17	99.75

Table 3: Monte Carlo estimates for % Type I error (a term is not selected when it actually belongs to the true model), % Type II error (a term is selected when it is not in the true model) and % F -statistic defined in Equation (13) using CLIC and CL-BIC criteria. Results are based on 500 Monte Carlo samples generated from Model 3 with $n = 10$ and $T = 10, 25$.

4 Analysis of the cancer cell growth data

The data analyzed in this section consist of counts of three cell types (cancer cells expressing GFP or mCherry, and fibroblasts) from 9 subsequent images taken at an 8-hour frequency over a period of 3 days using the Operetta high-content imager (Perkin Elmer). Each image was subsequently tiled using a 25×25 regular grid. GFP and mCherry cells express green (G) and red (R) fluorescent proteins respectively, while fibroblasts (F) do not express any fluorescent proteins. Cells were initially seeded at the ratio of 1:1:2. Fibroblasts are the predominant cell type in the tumor microenvironment and they are believed to affect tumor progression, partly due to interactions with and activation by cancer cells (Kalluri and Zeisberg, 2006).

Model selection and inference We applied our methodology to quantify the magnitude and direction of the growth for the considered cell types. To select the relevant terms in the intensity expression (1), we carry out model selection using the CL-BIC model selection criterion. In Table 4, we show estimated parameters for the full and the CL-BIC models, with bootstrap 95% confidence intervals in parenthesis. Figure 2 illustrates estimated spatio-temporal interactions between cell types using a directed graph. The solid and dashed arrows represent respectively the significant and not significant interactions between cell types at the 99% confidence level. Significant interactions coincides with parameters selected by CL-BIC.

The interactions within each cell type ($\hat{\beta}^{(c|c)}$, $c = R, G, F$) are significant, which is consistent with healthy growing cells. As anticipated, the effects $\hat{\beta}^{(c|c)}$ for the cancer cells are

larger than those for the slower growing fibroblasts. The validity of the estimated parameters is also supported by the similar sizes of the parameters for the green and red cancer cells. This is expected, since the red and green cancer cells are biologically identical except for the fluorescent protein they express. Interestingly, the size of the estimated effects within both types of cancer cells ($\hat{\beta}^{(c|c)}, c = R, G$) are larger than interactions between them ($\hat{\beta}^{(G|R)}$ and $\hat{\beta}^{(R|G)}$). This is not surprising, since $\hat{\beta}^{(c|c)}(c = R, G)$ reflects not only interaction between the same cell population, but also cell proliferation, suggesting that our model is sensitive enough to identify biologically relevant interactions.




Figure 2: Directed graph showing fitted spatio-temporal interactions between GFP cancer cells (G), mCherry cancer cells (R) and fibroblasts (F). The solid and dashed arrows represent respectively the significant and not significant interactions between cell types at the 99% confidence level. The thickness of the arrows is proportional to the size of the estimated effects.

Goodness-of-fit and one-step ahead prediction To illustrate the goodness-of-fit of the estimated model, we generate the number of cells for each type in each tile, $\hat{y}_{i,t}^{(c)}$, from the $\text{Pois}(\hat{\lambda}_{i,t}^{(c)})$ distribution for $t \leq 1$, where $\hat{\lambda}_{i,t}^{(c)}$ is computed using observations at time $t - 1$, with parameters estimated from the entire dataset. In Figure 3, we compare the observed and generated cell counts for GFP cancer cells (G) and mCherry cancer cells (R) and fibroblasts (F) across the entire image. The solid and dashed curves for all cell types are close, showing that the model fits the data reasonably well. As anticipated, the overall growth rate for the red and green cancer cells are found to be similar, and sensibly larger than the growth rate for fibroblasts.

To assess the prediction performance of our approach, we consider one-step-ahead fore-

$i =$	Full model		
	G	R	F
$\hat{\alpha}_i$	-0.99 (-1.12, -0.86)	-0.50 (-0.63, -0.37)	-0.26 (-0.33, -0.10)
$\hat{\beta}_{G i}$	1.23 (1.14, 1.30)	0.34 (0.26, 0.45)	0.12 (0.02, 0.25)
$\hat{\beta}_{R i}$	0.28 (0.22, 0.35)	1.09 (1.01, 1.16)	0.02 (-0.06, 0.08)
$\hat{\beta}_{F i}$	0.10 (0.02, 0.16)	0.02 (-0.06, 0.10)	0.92 (0.82, 0.98)
$i =$	CL-BIC model		
	G	R	F
$\hat{\alpha}_i$	-0.88 (-1.07, -0.74)	-0.49 (-0.67, -0.31)	-0.19 (-0.30, -0.02)
$\hat{\beta}_{G i}$	1.24 (1.09, 1.35)	0.35 (0.24, 0.48)	/
$\hat{\beta}_{R i}$	0.28 (0.19, 0.37)	1.09 (0.99, 1.18)	/
$\hat{\beta}_{F i}$	/	/	0.93 (0.78, 1.01)

Table 4: Estimated parameters for the full and the CL-BIC models based on the cancer cell growth data described in Section 4. Bootstrap 95% confidence intervals based on 50 bootstrap samples are given in parenthesis.

casting using parameters estimated from a moving window of five time points. In Figure 4, we show quantiles of observed cell counts against predicted counts per tile. The 95% upper and lower confidence bounds are computed non-parametrically as in Koenker (2005) by taking $\hat{F}_1^{-1}(\hat{F}_0(y_t^{(c)}) - 0.95)$ and $\hat{F}_1^{-1}(\hat{F}_0(y_t^{(c)}) + 0.95)$, where \hat{F}_0 and \hat{F}_1 are the empirical distributions of the observations and predictions at time t respectively. The identity line falls within the confidence bands in each plot, indicating a reasonable prediction performance.

5 Conclusion and final remarks

In this paper, we introduced a new conditional spatial autoregressive model and accompanying inference tools for multivariate spatio-temporal count data. The new methodology enables one to measure the overall cell growth rate in longitudinal experiments and spatio-temporal interactions with either homogeneous or heterogeneous cell populations. Inference is carried out by a new CL estimator which is computationally tractable and strikes a good balance between computational feasibility and statistical accuracy. Numerical findings from simulated and real data in Sections 3 and 4 confirm the validity of the proposed approach in terms of prediction, goodness-of-fit and estimation accuracy.

The proposed methodology can be used to answer many relevant questions in cancer




Figure 3: Goodness-of-fit of the estimated model. Observed (solid) and predicted (dashed) number of GFP cancer cells (G) and mCherry cancer cells (R) cancer cells and fibroblasts (F) for the entire image. Predicted cell counts for each cell type in each tile $\hat{y}_{i,t}^{(c)}$ is generated from the conditional Poisson model with intensity $\hat{\lambda}_{i,t}^{(c)}$ defined in Equation (1) and (2), where the coefficients $\hat{\beta}^{(c|c')}$ are estimated from the entire dataset.

cell biology. Examples include how other cell types (alone or combined) in the tumor microenvironment, such as macrophages, tumor-infiltration lymphocytes and other immune cells, affect cancer cell growth (Ungefroren et al., 2011). It is also possible to study the effect of drugs on these interactions as well as how different genes and associated pathways in certain cell types are involved in inter-cellular cross talk by genetically modifying the cell type(s) in question before mixing the cells together. Our approach is also applicable in other kinds of studies where local spatial cell-cell interactions are believed to affect cell growth such as studies of neurodegenerative diseases (Garden and La Spada, 2012) and wound healing/tissue re-generation (Leoni et al., 2015). In addition to evaluating cell growth, our approach can also be used to study transitions between cellular phenotypes upon interaction with other cell types, provided that phenotypes can be distinguished based on the image data.

Our methods offer several practical advantages to researchers interested in analysing multivariate count data on heterogeneous cell populations. First, the conditional Poisson model does not require tracking individual cells across time, a process that is often difficult to automate due to cell movement, morphology changes at subsequent time points, and additional complications related to storage of large data files. Second, we are able to quantify local spatio-temporal interactions between different cell populations from a very simple experimental set-up where the different cell populations are grown together in a single experimental condition (co-culture). An alternative solely experimentally-based strategy




Figure 4: QQ-plots for cell growth, comparing observed (horizontal axis) and one-time ahead predicted (vertical axis) cell counts per tile on the entire image at times $t = 6, 7, 8$ for GFP cancer cells (G), mCherry cancer cells (R) and fibroblasts (F). One-time ahead predictions are based on the model fitted using a moving window of five time points.

would require monitoring the different cell types alone and together at different cell densities (number of cells per condition) in order to make inferences in terms of potential interactions. However, such an approach would give no possibility of evaluating the spatial relations in the co-culture conditions and would still restrict the number of simultaneously tested cell types to two.

In the future, we foresee several useful extensions of the current methodology, possibly enabling the treatment of more complex experimental settings. Complex experiments involving a large number of cell populations, n_C , would imply an over-parametrized model. Clearly, this large number of parameters would be detrimental to both statistical accuracy and reliable optimization of the CL objective function $cl(\theta)$ (5). To address these issues,

we plan to explore a penalized composite likelihood of form $c\ell(\theta) - \text{pen}_\lambda(\theta)$, where $\text{pen}(\theta)$ is a nonnegative sparsity-inducing penalty function. For example, in a different CL setting Bradic et al. (2011) consider the L_1 -type penalty $\text{pen}(\theta) = \lambda \sum |\theta|$, $\lambda > 0$. Finally, for certain experiments, it would be desirable to modify the statistics in (2) to include additional information on cell growth such as the distance between heterogeneous cells, and covariates describing cell morphology.

6 Acknowledgements

The authors wish to acknowledge support from the Australian National Health and Medical Research Council grants 1049561, 1064987 and 1069024 to Frédéric Hollande. Christina Mølck is supported by the Danish Cancer Society.

References

- A. Amatya and H. Demirtas. Poisor: An R package for generation of multivariate data with Poisson and Normal marginals. *Communications in Statistics-Simulation and Computation*, (just-accepted), 2015.
- J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 192–236, 1974.
- F. D. Bowman. Spatiotemporal models for region of interest analyses of functional neuroimaging data. *Journal of the American Statistical Association*, 102(478):442–453, 2007.
- J. Bradic, J. Fan, and W. Wang. Penalized composite quasi-likelihood for ultrahigh dimensional variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(3):325–349, 2011.
- R. Chen, O. Hyrien, M. Noble, and M. Mayer-Pröschel. A composite likelihood approach to the analysis of longitudinal clonal data on multitype cellular systems under an age-dependent branching process. *Biostatistics*, 12(1):173–191, 2011.

- G. Claeskens, N. L. Hjort, et al. *Model selection and model averaging*, volume 330. Cambridge University Press Cambridge, 2008.
- N. Cressie. *Statistics for spatial data*. John Wiley & Sons, 2015.
- K. Fokianos and B. Kedem. Partial likelihood inference for time series following generalized linear models. *Journal of Time Series Analysis*, 25(2):173–197, 2004.
- K. Fokianos and D. Tjøstheim. Log-linear Poisson autoregression. *Journal of Multivariate Analysis*, 102(3):563–578, 2011.
- X. Gao and P. X.-K. Song. Composite likelihood Bayesian information criteria for model selection in high-dimensional data. *Journal of the American Statistical Association*, 105(492):1531–1540, 2010.
- G. A. Garden and A. R. La Spada. Intercellular (mis) communication in neurodegenerative disease. *Neuron*, 73(5):886–901, 2012.
- A. E. Gelfand, H.-J. Kim, C. Sirmans, and S. Banerjee. Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association*, 98(462):387–396, 2003.
- C. Heyde. On combining quasi-likelihood estimating functions. *Stochastic Processes and their Applications*, 25:281–287, 1987.
- O. Hyrien. Pseudo-likelihood estimation for discretely observed multitype Bellman–Harris branching processes. *Journal of statistical planning and inference*, 137(4):1375–1388, 2007.
- O. Hyrien, M. Mayer-Pröschel, M. Noble, and A. Yakovlev. A stochastic model to analyze clonal data on multi-type cell populations. *Biometrics*, 61(1):199–207, 2005.
- R. Kalluri and M. Zeisberg. Fibroblasts in cancer. *Nature Reviews Cancer*, 6(5):392–401, 2006.
- H. Kang, H. Ombao, C. Linkletter, N. Long, and D. Badre. Spatio-spectral mixed-effects model for functional magnetic resonance imaging data. *Journal of the American Statistical Association*, 107(498):568–577, 2012.

- D. Karlis and L. Meligkotsidou. Multivariate Poisson regression with covariance structure. *Statistics and Computing*, 15(4):255–265, 2005.
- R. Koenker. *Quantile regression*. Number 38. Cambridge university press, 2005.
- G. Leoni, P. Neumann, R. Sumagin, T. Denning, and A. Nusrat. Wound repair: role of immune–epithelial interactions. *Mucosal immunology*, 8(5):959–968, 2015.
- B. G. Lindsay. Composite likelihood methods. *Contemporary mathematics*, 80(1):221–39, 1988.
- H. Ungefroren, S. Sebens, D. Seidl, H. Lehnert, and R. Hass. Interaction of tumor cells with the microenvironment. *Cell Communication and Signaling*, 9(1):1, 2011.
- C. Varin and P. Vidoni. A note on composite likelihood inference and model selection. *Biometrika*, 92(3):519–528, 2005.
- C. Varin, N. M. Reid, and D. Firth. An overview of composite likelihood methods. *Statistica Sinica*, 21(1):5–42, 2011.
- Y. Zheng and J. Zhu. Markov chain Monte Carlo for a spatial-temporal autologistic regression model. *Journal of Computational and Graphical Statistics*, 2012.
- J. Zhu, H.-C. Huang, and J. Wu. Modeling spatial-temporal binary data using Markov random fields. *Journal of Agricultural, Biological, and Environmental Statistics*, 10(2):212–225, 2005.
- J. Zhu, J. G. Rasmussen, J. Møller, B. H. Aukema, and K. F. Raffa. Spatial-temporal modeling of forest gaps generated by colonization from below-and above-ground bark beetle species. *Journal of the American Statistical Association*, 103(481):162–177, 2008.